

LA-UR- 09-05118

Approved for public release;  
distribution is unlimited.

Title: Kerberized Network File System for Clusters

Author(s): Ian Burns  
Christopher Hoffman  
Paige Ashlynn

Intended for: Electronic/World-wide Web  
Academic Distribution



Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

# Abstracts

---

## Kerberized Network File System for Clusters

Paige Ashlynn, UNM

Ian Burns, NMT

Christopher Hoffman, MTU

**Abstract:** With constantly looming cyber-security threats, protecting valuable data has become a very important issue. The implementation of security protocols should not compromise ease of use or performance. The Kerberos protocol provides a high level of security while minimizing overhead. A central Kerberos server needs to be able to provide authentication for a variety of services distributed over a number of connected networks. Though clustered environments traditionally have not required internal security, the landscape is changing rapidly.

It is important for any authorized person to be able to access their data from whatever computer they must use for their work. This could be a simple desktop workstation, or a large supercomputer. There needs to be a single, secure method of accomplishing this sharing for all environments. Kerberized NFS can be used to address this need for data mobility in a secure manner. However, the performance impact that Kerberos will have on NFS in a clustered setting is still largely unknown. Factors such as level of security and different types of encryption affect performance and usability greatly.

We will evaluate these impacts and make a general recommendation for suitable security levels and feasibility for possible deployments in current and future LANL systems.

---

# Kerberized Network File System for Clusters

Presented on 08-03-2009 by:

Chris Hoffman

Ian Burns

Paige Ashlynn

Instructor: Andree Jacobson

Mentor: David Kennel



UNCLASSIFIED



Operated by Los Alamos National Security, LLC for NNSA



# Introduction

---

## Kerberized NFSv4 in an HPC Cluster

- Motivation
- Advantages & Disadvantages of Kerberos
- Changes in NFSv4
- Behaviour
- Performance

# Impetus

---

## The Problem:

- Secure, light-weight, remote data access from nodes on supercomputer clusters.

## The Solution:

- Kerberos and NFSv4 are designed to accomplish this task on enterprise networks.

# Advantages of Kerberos Authentication

---

- Single Sign-On
- Password Security
- Verified Clients
- Scalable
- Integration

# Disadvantages of Kerberos Authentication

---

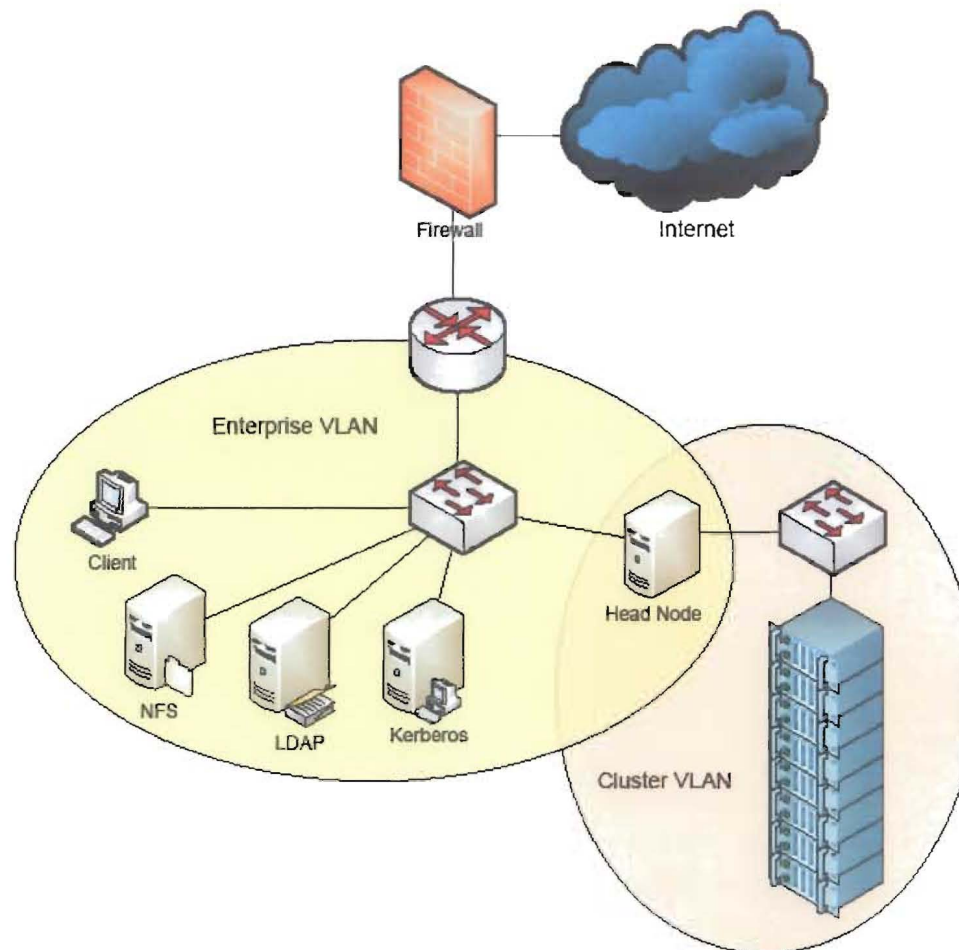
- Single Point of Compromise
- Requires Application Support

# Changes in NFSv4

---

- Native Support for Authentication Protocols
- Kerberos Integration Options:
  - KRB5 = Authentication
  - KRB5i = Authentication & Integrity
  - KRB5p = Authentication, Integrity, & Encryption
- Non-Unix Compatibility
- Virtual Filesystem

# Test Environment



# Behaviour in a Clustered Environment

---

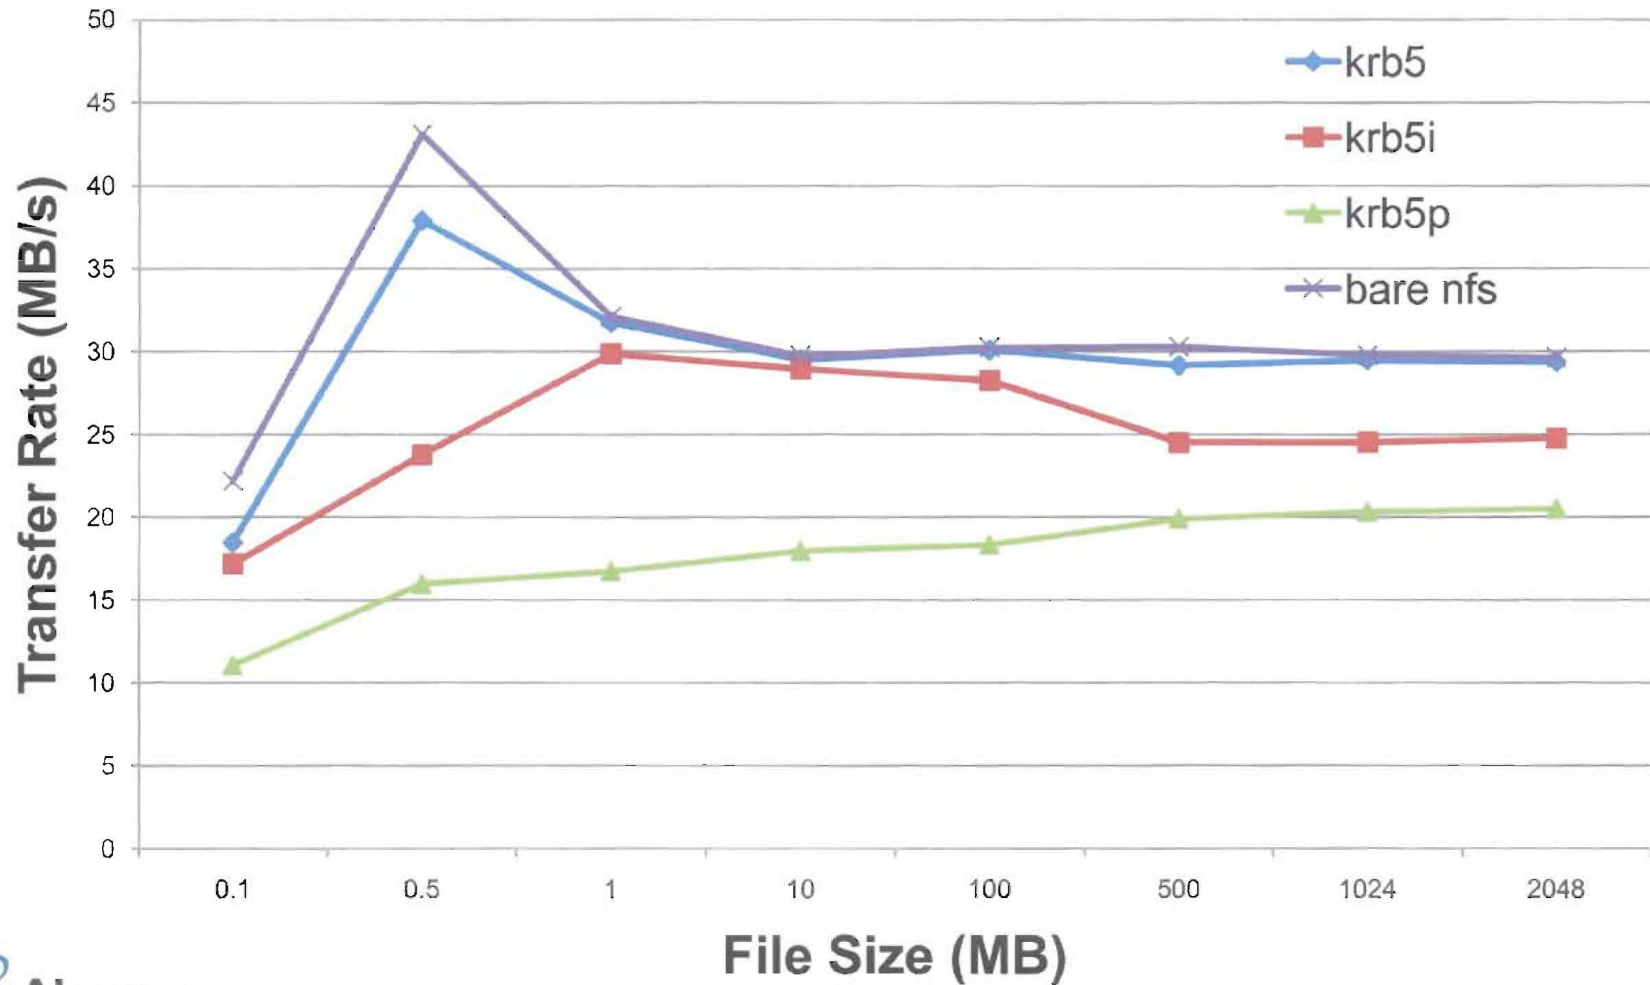
- Network Address Translation
- Addressless ticketing
- Torque Job Scheduler

# Performance Test

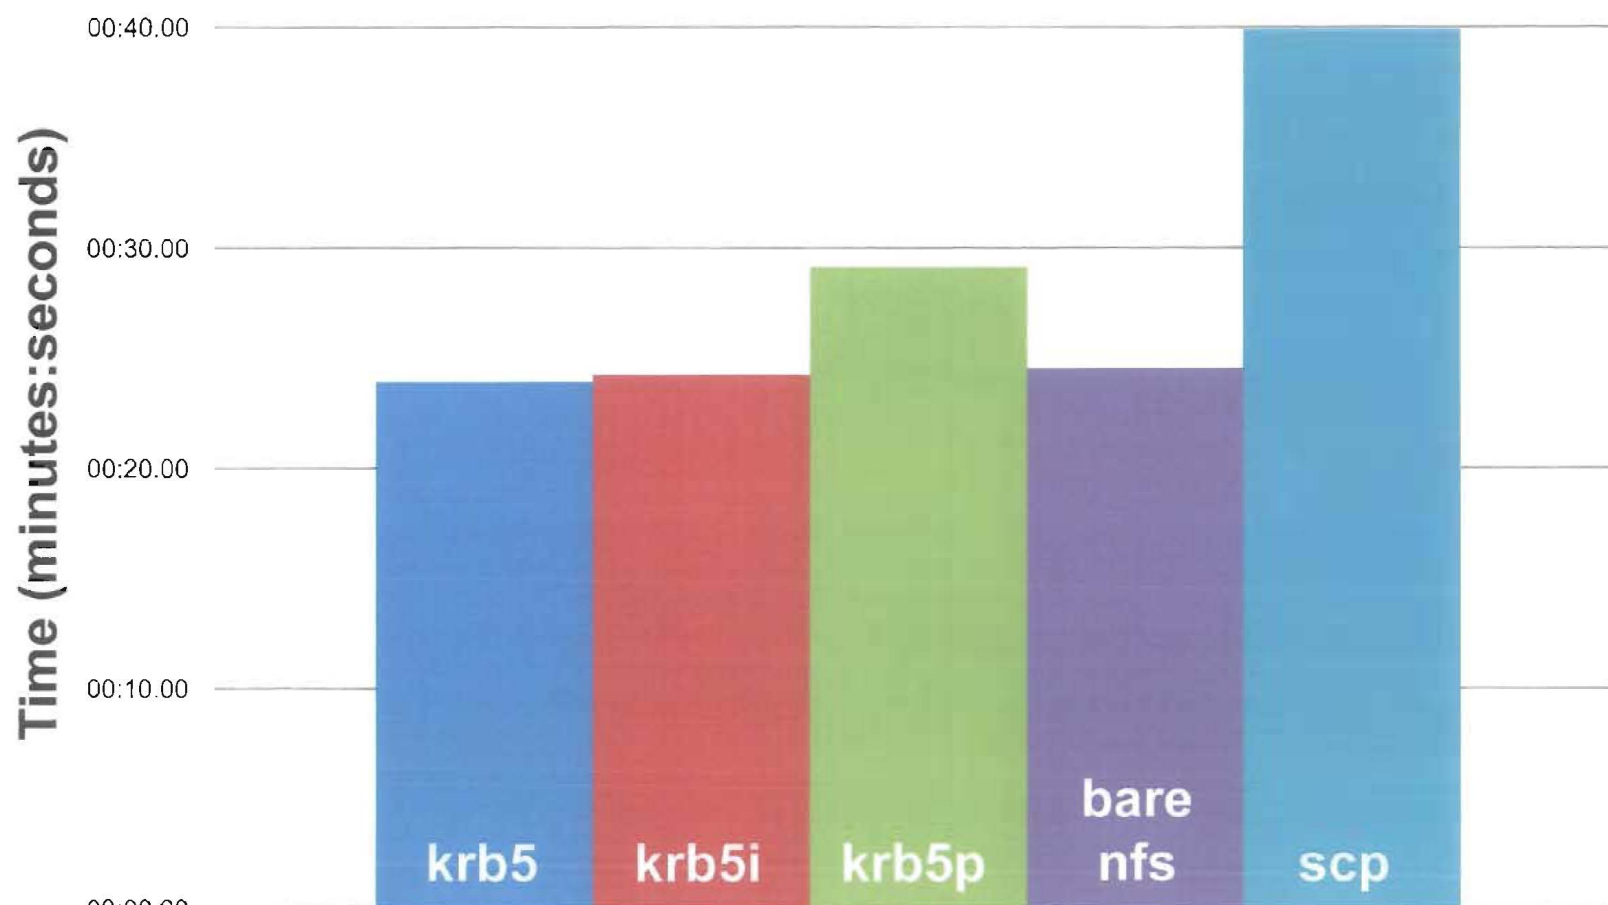
---

- Levels of Security
- Large vs. Small Files
- SCP vs. Bare NFSv4 vs. Kerberized NFSv4  
over Gigabit Ethernet

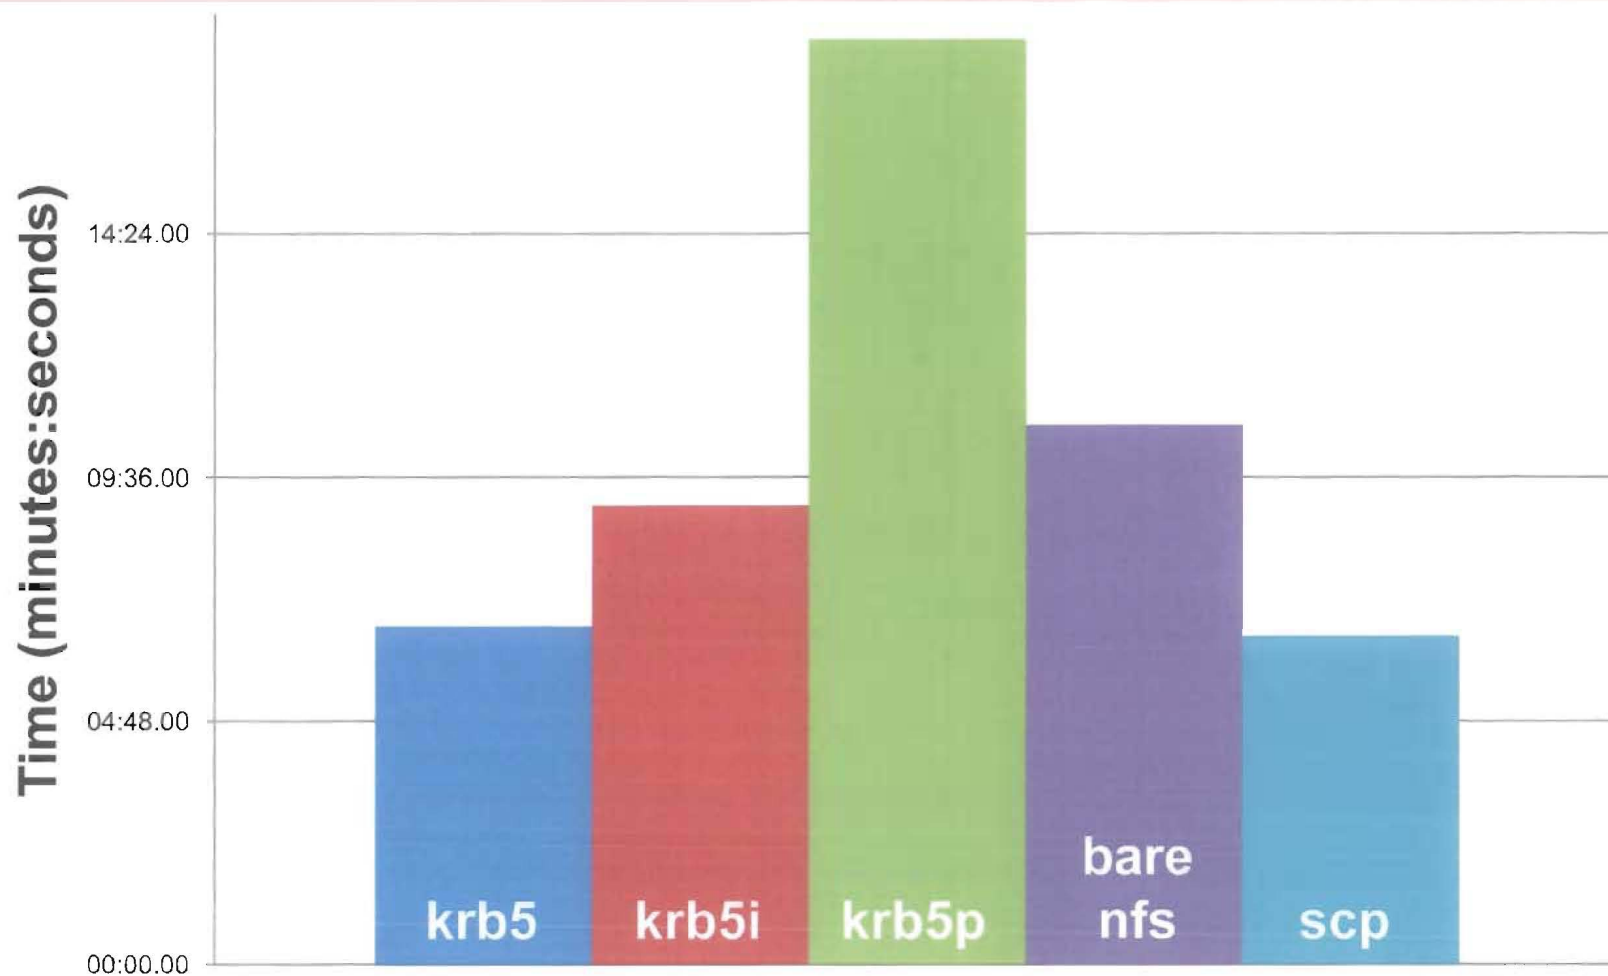
# Average File Transfer Rate



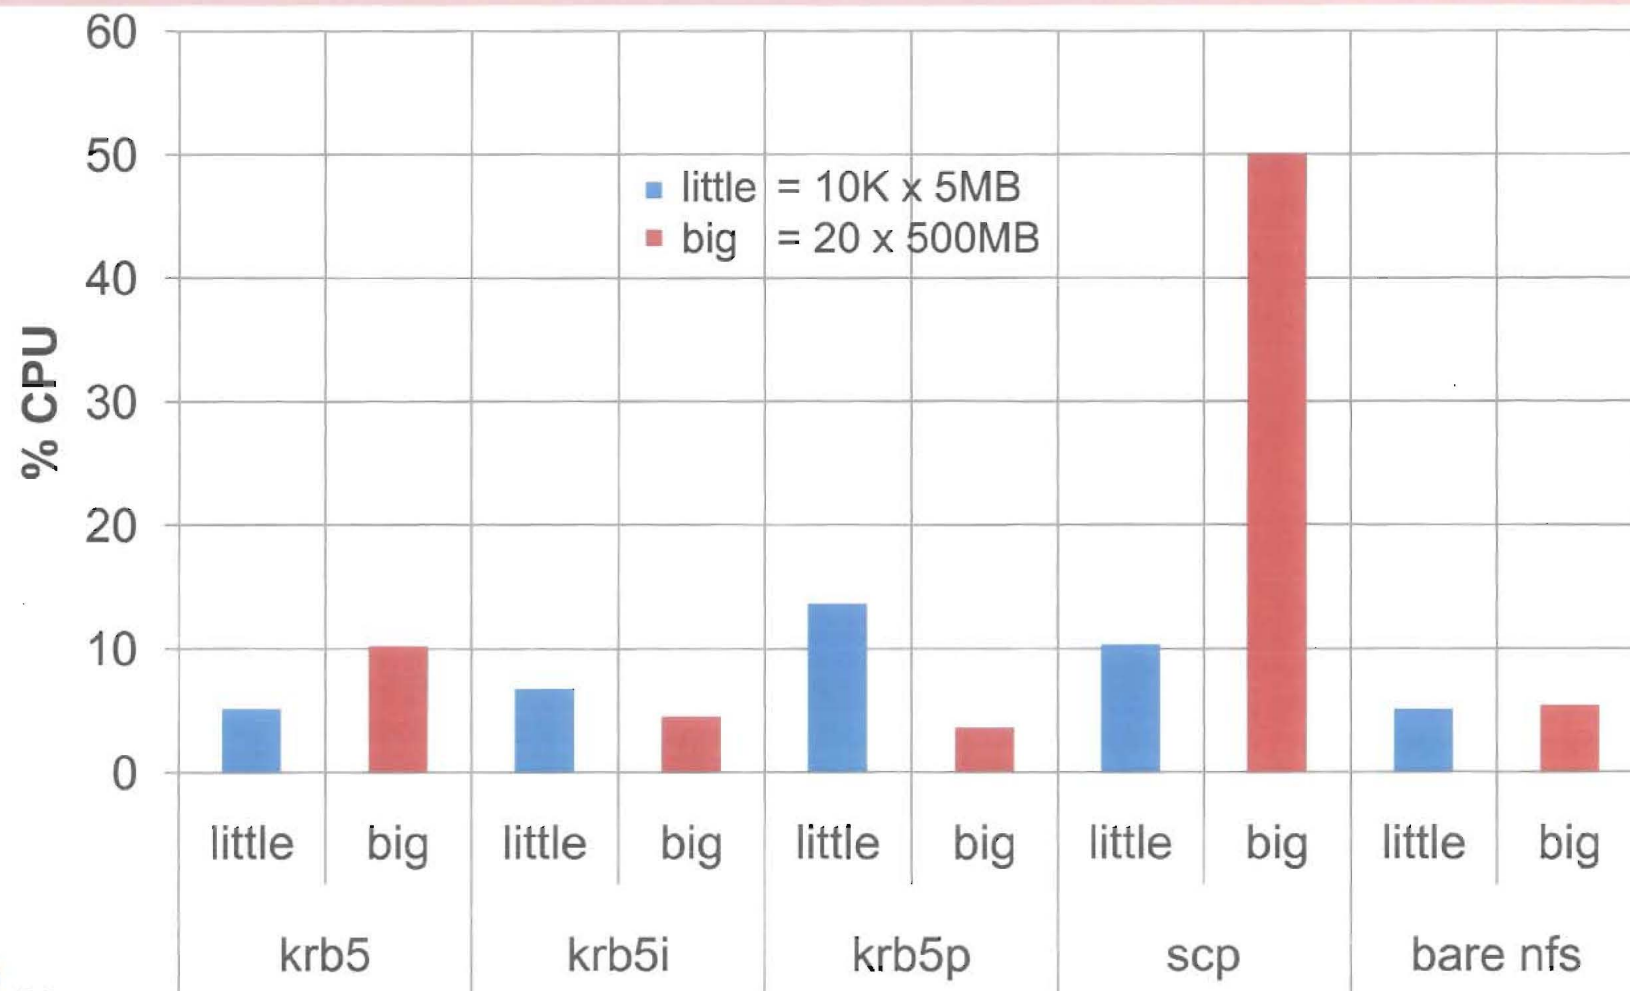
# Average Time to Copy 10,000 5-kB Files



# Average Time to Copy 20 500-MB Files



# CPU Utilization by Protocol



# Summary

---

## Kerberized NFSv4 in an HPC Cluster

- Motivation
- Advantages & Disadvantages of Kerberos
- Changes in NFSv4
- Behaviour
- Performance

# Questions?

## Answers!

